DOCUMENT RESUME

ED 081 282 FL 004 579

| | |
|---|---|
| AUTHOR | Schank, Roger C.; Wilks, Yorick |
| TITLE | The Goals of Linguistic Theory Revisited. |
| INSTITUTION | Stanford Univ., Calif. Dept. of Computer Science. |
| SPONS AGENCY | Department of Defense, Washington, D.C. Advanced Research Projects Agency. |
| REPORT NO | M-AIM-202; STAN-CS-73-368 |
| PUB DATE | May 73 |
| NOTE | 43p. |
| EDRS PRICE | MF-$0.65 HC-$3.29 |
| DESCRIPTORS | Algorithms; Computational Linguistics; Computer Programs; Deep Structure; Grammar; Information Processing; Language Research; *Language Usage; *Linguistic Competence; *Linguistic Performance; *Linguistic Theory; Logic; Memory; Prediction; *Semantics; Sentence Structure; Surface Structure; Transformation Generative Grammar; Transformation Theory (Language) |

ABSTRACT

There is a need for a new kind of linguistic theory which, while being concerned with both generation and analysis, must include the roles of memory, non-linguistic knowledge, and inference. The role of logic is diminished according to such a theory because inference has no real logical content. Meaning must be studied with respect to the actual usage of the speaker. Linguistics must encompass a procedure for making explicit the information that is implicit in certain sentences, and it must allow for the possibility of mistakes and mistaken inferences in speech. Linguistic theory must deal with the way humans use language and the facts of language use. (VM)

FILMED FROM BEST AVAILABLE COPY

ED 081282

STANFORD ARTIFICIAL INTELLIGENCE LABORATORY
MEMO AIM-202

COMPUTER SCIENCE DEPARTMENT
REPORT NO. CS-368

MAY 1973

# THE GOALS OF LINGUISTIC THEORY REVISITED

by

Roger C. Schank
Yorick Wilks

U S DEPARTMENT OF HEALTH, EDUCATION & WELFARE
NATIONAL INSTITUTE OF EDUCATION
THIS DOCUMENT HAS BEEN REPRODUCED EXACTLY AS RECEIVED FROM THE PERSON OR ORGANIZATION ORIGINATING IT. POINTS OF VIEW OR OPINIONS STATED DO NOT NECESSARILY REPRESENT OFFICIAL NATIONAL INSTITUTE OF EDUCATION POSITION OR POLICY

ABSTRACT: We examine the original goals of generative linguistic theory. We suggest that these goals were well defined but misguided with respect to their avoidance the problem of modelling performance. With developments such as Generative Semantics, it is no longer clear that the goals are clearly defined. We argue that it is vital for linguistics to concern itself with the procedures that humans use in language. We then introduce a number of basic human competencies, in the field of language understanding, understanding in context and the use of inferential information, and argue that the modelling of these aspects of language understanding requires procedures of a sort that cannot be easily accomodated within the dominant paradigm. In particular, we argue that the procedures that will be required in these cases ought to be linguistic, and that the simple-minded importation of techniques from logic may create a linguistics in which there cannot be procedures of the required sort.

This research was supported by the Advanced Research Projects Agency of the Department of Defense under Contract SD-183.

The views and conclusions contained in this document are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, of the Advanced Research Projects Agency or the U.S. Government.

Reproduced in the USA. Available from the National Technical Information Service, Springfield, Virginia 22151.

FL004 579

I. What is a linguistic theory best considered to be a theory of? Let us begin by asking what the original goal of modern linguistics was when Chomsky began directing its course. We can start profitably from Chomsky's view that a linguistic theory is a "theory of a language L", and that such a theory is scientific in the normal sense of that word:

"A grammar of the language L is essentially a theory of L. Any scientific theory is based on a finite number of observations, and it seeks to relate the observed phenomena and to predict new phenomena by constructing general laws in terms of hypothetical constructs. . . Similarly a grammar of English is based on a finite corpus of utterances (ob- servations), and it will contain certain grammatical rules (laws) stated in terms of the particular phonemes, phrases, etc. , of English (hypothetical constructs). These rules express structural relations among the sentences of the corpus (predictions). "

Fairly straightforward considerations tell against this way of looking at transformational grammars. For, in the case of scientific theories cast in standard hypothetico-deductive form, there is a well-understood notion of what it is to disconfirm a particular theory. There are difficulties about making this notion of disconfirmation precise; nonetheless there is general agreement about both its form and its importance. But Chomsky, in the

quotation above, has formulated the theory of transformational grammar so as to rule out the possibility of disconfirmation. When describing what an 'utterance' is for the purpose of inclusion in a 'corpus', Chomsky makes clear that he is not going to include what appear to be utterances, but which are strings containing 'grammatical mistakes' [1] . The notion of 'grammatical mistake' is defined with respect to the grammar in question, so that there can never be a rejected grammar.

What we have shown non-predictive, and therefore not scientific in the desired sense, are what Chomsky calls 'competence theories', so perhaps we can persist a little with the question of what a linguistic competence theory is a theory of; given that it cannot, by definition for Chomsky, be brought into direct contact with human behavior (for that is the scope of 'performance'), and is not intended to be a brain model either (. . . "the deeper absurdity of regarding the system of generative rules as a point-by-point model for the actual construction of a sentence by a speaker" [2] , where we are taking the "point-by-point" phrase to mean something that could be no other than a brain model.

The fact of the matter is that linguistic theory, apart from its traditional classificatory and comparative concerns, just will not fit into any acceptable form for being "scientific". The classificatory concerns were and are scientific, in exactly the way

that Linnaeus' plant classification was scientific in its time. But linguistics cannot be forced into some other paradigm of science, at the present time, such as that of the hypothetico-deductive theory. If a linguistic theory is not a scientific theory in the normal sense, and is not a theory of human behavior or of the brain, then what is it? What could be meant by the only reply left open, "well, then, it is a non-scientific theory of a language. " We would maintain that Chomsky's theories have always been, in a sense, theories of productive mechanisms or algorithms. A perceptive remark of Putnam's will illustrate the point:

". . . . the reader. . . may go through a work like Chomsky's Syntactic Structures carefully, and note that at no place is the assumption employed that the corpus of utterances studied by the linguist was produced by a conscious organism". [12]

We can bring this observation up to date by quoting a more recent statement of Chomsky's on the nature and role of a grammar (that is to say a theory of competence):

". . . . by a generative grammar I mean simply a system of rules that in some explicit and well defined sense assigns structural descriptions to sentences. . . . . The term 'generate' is familiar in the sense intended here in logic, particularly in Post's theory of combinatorial systems". [2]

These quotations are only to remind the reader that the

algorithmic, or device oriented, way of talking about linguistic theories is already familar, and is utterly different from model, psychological entity, scientific theory, or reality in the brain, modes of talk.

Chomsky's own comparison with Post's logic makes the point precise: Post's logic was productive, or generative, in the simple old-fachioned sense of those words. Such a logic consisted of rules, written with short left-hand sides and long right-hand ones, and produced progressively longer objects called theorems, and so, in the linguistic case, correct sentences. Because of the progressive obfuscation of the term "generate" in recent years. it is important to make this simple point clear: that Chomskyan generative linguistics began as a system of rules for producing sentences.

Chomsky's original self-imposed task then, was the description of a mechanism that would generate all and only the language strings satisfying some criterion of correctness. That remains the fundamental description of what Chomsky was aiming at, even though it is now called "weak generative capacity, " and the criterion of correctness itself has wobbled a bit over the years.

There seems to be a continuing confusion in current linguistics on this point, in that, in their eagerness to disclaim any intention to model the mind, brain, or other processes of an actual speaker, some linguists have gone too far and disavowed the

original notion of sentence production as well. The task of the linguist is then thought to be no more than assigning descriptions to individual sentences, though by methods which must remain wholly mysterious if he has already rejected all actual analytic or productive algorithms.

The weakest form of this doctrine, if we turn now to current linguistics, is the informal use of words like "blueprint" to describe the function of competence: a grammar is then a blueprint "referred to in the construction of sentences". This way of speaking captures the worst of all possible worlds, in that it lacks even the definite falseness of those who, wrongly, as one of us has argues elsewhere, [20] speak of models in this context. To speak of a model is to commit oneself, as it is to a lesser degree to speak of a theory, and, as we will argue below at length, to to speak of the construction of precise bodies of analytic or generative rules commits oneself in the most positive way at this stage of the development the discipline of linguistics. But the talk of "blueprints" commits one to nothing, and seems to us to attach itself to no precise activity at all.

A good place to look, if we ask what is the goal of current linguistic theory, is to the school of generative semanticists. Lakoff [7] describes their enterprise as follows:

"Generative semantics [GS] claims that the underlying

grammatical structure of a sentence is the logical form of that sentence, and consequently that the rules relating logical form to surface form are exactly the rules of grammar. "

It may well be the case that certain of the rules to which Lakoff has drawn attention in his paper do have a part to play in any general language-to-logic translation AND in any reasonably general grammar, of whatever sort. But that is a far cry, of course, from the burden of proof required by the "exactly" in the last quotation. If it is replied that the quotation expresses only a conjecture, then it seems a false one, since it is not hard to find ----for two such prima facie different tasks as grammatical production, and translation of language to logic---- examples of rules that will certainly function in one enterprise and equally certainly not in the other. We do not believe, for example, that the grammaticality of sentences containing "possibly" can require a rule relating that word to some primitive symbol expressing the concept of certainty. Yet translation of such sentences into modal logic will require some such rule (or the complement of it, where "certain" replaces "possible" mutatis mutandis). Surely Lakoff's conjecture-assertion about rule identity excludes this possibility?

With GS, as with all such theses, there are two ways of looking at it: one is to take the words as meaning what they appear to mean; the other is to assume that they mean something quite

different. The first approach gives us the TRANSLATION view, or the CONSEQUENCE view, depending on how we take the word "relating" in the last quotation. The second approach would give the RENAMING view: on which, when Lakoff speaks of logical form, he does not mean any standard sense of the phrase, but rather some linguistic structure, either familiar or one of his own devising. In either case, on the renaming view, GS would not really be about logic at all, and disputes about the GS thesis would be wholly an internal matter for linguistics. When Chomsky [3] writes of GS as "notational variant" of his own work, he is taking what we call the renaming view.

The consequence view is the most obvious possibility, namely that the "relates" is by inference, valid or otherwise, and that the well-formedness of sentences is settled by whether or not they can be inferred from logical forms. Much of the evidence for this assumption is circumstantial because Lakoff rarely discusses GS in general terms. But it is reinforced by his introduction of rules of inference with "It is clear that there is more to representing meanings than simply providing logical forms of sentences"[ibid]. That quotation seems to us to rule out the translation view: that logical forms are the meaning, or "backbone", or sentences and can be related to them by mere rules of translation. The translation view also becomes less plausible when one remembers how much of Lakoff's work is about inference: if GS were really about translation into

logical form, then inference would have no place at all in a discussion of natural logic. So then, the consequence view must be Lakoff's view, if he has a firm view. Two clear and simple considerations tell against it:

(1) There is no clear notion available of inference that goes from logical forms to sentences. Rules that cross the logical form-sentence boundary are rules of translation.

(2) There is the problem of "reverse direction": how could we analyse sentences with reverse inference rules to produce falsehoods, as in "if this is not colored then it is not red." What possible interpretaion could we attach to such a procedure in the context of GS?

This last is the key point for the underlying question we are discussing. The doubtful word in the definition of GS is "relate": it has all the directional ambiguity of "generate". However, in nailing his colors to the mast of logic, and to a logic of inferential relations at that, Lakoff has also, perhaps unknowingly, committed himself to sentence production as the real goal of his linguistic theory since, as we have shown, inference rules simply cannot be reversed to yield analysis routines. Yet, nonetheless as we shall argue in the next section, informal sentence-by-sentence analysis is what generative semanticists actually do when they do linguistics, and this utter confusion

between proclaimed and actual goals is one of the most vulnerable spots in the armor of modern linguistics.

Furthermore, although linguists insist that what they are doing is a directionless relating of sound and meaning, it is, in fact, extremely difficult to relate sounds to meanings (in that order) using transformational grammar. If deletion transformations are allowed, the task is simply impossible mathematically. If not, the nature of many-to-one mapping makes the problem of finding a path from sound to meaning to be so much one of trial and error as to be useless as a basis for a precise theory.

Thus, we claim, that generative linguists are working on a task to which they cannot provide any reasonable solution , given their initial assumptions.

II. In spite of what we argued in the last section about the theoretical problem of reversiing TG and GS rules, it is the proclaimed goal of most current linguistic theory to make explicit the underlying competence of speakers of a language by assigning an interpretation and structural description to grammatical sentences.

Let us look at the process by which a generative linguist actually decides what the interpretation and structural description of a given sentence is. First the linguist considers the sentence, and decides on its meaning in his own mind. He then

writes, according to the rules for creating structural diagrams, a structural description for the meaning of that sentence. If he is now to go further with this sentence he will address himself principally to two issues. First, what would be the form of the explicit rules that would transform this structural diagram into a surface structure for this sentence? Second, how do the rules that would have to be created to do the first task conflict with what the same grammatical rules and structural diagrams have been previously understood to be? That is, must the rules for writing structural diagrams be modified, or must those used for assigning surface structure be emended in order to have a consistent theory?

If this is, in fact, a fair description of the procedures of generative linguists, it is interesting to inquire what is actually being done, as opposed to what the stated aims of generative linguistics are. A basic premise of generative theory is that rules are to be given that "relate" deep structures to surface structures without regard to the direction of this relationship. Do generative linguists actually provide such rules?

The answer is that they most certainly do not. The actual grammar rules to be found in any generative grammar are uniformly one-directional. They proceded from deep structures to surface structures and that is all. We may assume that a generative grammar

is not really intended to relate surface structures to deep structures since no hint of a possible procedure for doing this has ever been given by a linguistic theorist working on generative grammar. And, as we argued in section 1 , this omission is no accident because , on a "consequence" interpretation of GS, such rules could not be given.

But, in fact, generative grammarians do not actually write grammars which map deep structure into surface structures either. Although some few attempts have been made to write such a grammar, there is actually no complete grammar available to someone who might want to use it.

So what do generative grammarians do after all? What they actually do is work on an informal theory of semantic or syntactic representation , and discuss the problems in volved in relating(in one direction) this representation to what is considered a grammatical structure.

So, since generative grammarians actually work on producing semantic and syntactic representations of surface sentences, they are in fact doing analysis rather than generation, although they are making no attempt to specify the procedure by which they do such an analysis. Thus generative grammarians write structural diagrams for sentences, yet they make no claim to know how they do this analysis. Yet the fact that the 'deep structures' for sentences are

arrived at analytically implies that the question of analysis should be central for generative linguists, even if it is only dealt with informally.

This leads to the question of what an analytic linguistic theory would have in common with a generative linguistic theory should both exist. This problem was tackled by computational linguists in the early days of mechanical translation research. Often the approach was to use precisely the same rules for analysis as for generation, by simply reversing them. Others tried to have separate systems for each process with a system of transfer rules to take the output of the analytic routine into the base of the generative routine. Clearly linguists do this informally when they write generative grammars. That is, they first analyze into the semantic representation ( S.R. ), write it down and use it to think about generative grammars. If linguists do this, and more importantly if speakers of a language do this, then the problem of deciding what is necessary for an S.R. for a generative grammar, is at least partially dependent on the speaker's analytic procedure. That is, if certain things are necessary for an S.R. in order for it to be an adequate analysis, then it is reasonable to assume that this analytic base could also be used as an effective generative base, and that it should be used as such. In other words, it is possible to decide the adequacy of a particular base for a generative

grammar on the adequacy of that base as an output from an analytic procedure.

Linguistic theory should be concerned with both analysis and generation. Any really adequate theory should provide a base component that is capable of not only analyzing and generating but also connecting with a memory that could provide the input to the generative procedure and operate on the output of the analytic one .

In providing a detailed, programmable base that is useful for analysis and generation, problems arise that shed light on the question of the adequacy of current generative theories. For example, a good analysis of a sentence often contains references to items that are not explicitly present in the surface structure of that sentence, which would indicate that representations containing more information are to be preferred over less expressive representations. (This will be discussed further in section 7. )

If an analytic procedure is to add additional information it must contain predictive mechanisms so that it is possible to know when that information is needed. Thus, the S.R. used for analysis must be explicitly defined so that it can direct the analysis by looking at either the sentence or its memory for information that is predicted from the context by the formal base structure.

People who work with computers know that an analysis used in a

conversation program that cannot make the likely inferences will cause the program to function unintelligently. Furthermore, humans who cannot make appropriate inferences do rather a bad job of understanding things said to them. The fact is that inferences are an important part of the linguistic process and linguistic theory has to deal with them. However, while some linguists would agree with this point, few would allow their theory to make inferences that are only possibly true. The fact that inferences can be wrong should not be a deterrent to making them within the context of a linguistic model. People misinfer all the time; they correct themselves when they are wrong, but that is the nature of conversation. It is absolutely necessary that linguistic theory deal with this ability as part of competence, and not relegate it to the Siberia of performance. We shall argue now that, although making inferences is essential to an adequate linguistic theory, the current generative paradigm simply cannot accomodate this possibility in a serious way and that is one of the main things wrong with it.

III. In this paper, our main aim is to set out, in sketch form, what we feel the goals of an adequate linguistic theory should be at the present time. Here we would argue that it is a new sort of linguistic theory we need, and that it is a mistaken act of desperation to look, as the generative semanticists do, to logic to provide what linguistics cannot.

Let us make this point by looking close at what Lakoff means by a natural logic.

Lakoff writes [7]:

"(iv) We want a logic in which all the concepts expressible in natural language can be expressed unambiguously, that is, in which all non-synonymous sentences. . . . have different logical forms.

(v) We want a logic which is capable of accounting for all correct inferences made in natural language and which rules out incorrect ones. We will call any logic meeting the goals (above) a "natural logic"".

Again [ibid. ]

"In natural logic. . . . logical equivalences could not just be arbitrarily set down; rather they would be just those necessary to characterize the notion "valid inference" for natural language arguments".

And again [ibid. ]:

"Natural logic, taken together with linguistics, is the empirical study of the nature of human language and human reasoning".

This all sounds a very nice idea, and generally a good thing, but what does it really come to? These quotations, for example, taken together, express a curious ambivalence towards formal logic that runs right through that paper[7]. He writes of a natural logic in terms of the general study of human reasoning, but the fact

is that most real human reasoning is of a sort that is of interest to no one but psychologists, and sometimes psychiatrists. Real people argue much of the time along the lines of "That man has a squint, therefore he probably wants to mug me". And, of course, sometimes they are right in such inferences. The notion of inference, as such, has no real logical content: inferences are just the inferences that people actually make. Philosophers from Moore [11] to Lakoff [ibid. ], have criticised the basic connective "⊃", of material implication, in the Propositional Calculus, on the grounds that it in no way expressed the natural usage of "if. . . then" in ordinary language, because it allows any statement to imply any other, as as long as the first is not true while the second is false. In the Propositional Calculus one could truly say that "The Apollo space craft is nearing the moon" implies "I have a head-ache coming on", if indeed I do. But, more seriously, Lakoff also refers, in the passages quoted, to "valid" and "correct" inference when setting out what a a natural logic is to be. "Valid" is a reasonably well-understood term and covers such inferences as "all f's are g and all g's are l, therefore all f's are l", as well as those like "John is a younger son, therefore John has a brother".

We can easily construct a sense of "correct" inference, too, different from that of "valid inference" but still of interest to logic. For example, and to use an old logical favorite, we can

infer from "This is a creature with a heart" that "This is a creature with a liver". We can do this because the missing premise is universally true, since all creatures with hearts do as a matter of fact have livers, though this inference does not depend on the meanings of words as does the "younger son" case. But such inferences will be correct in that they will (while the world stays roughly the same as now) always lead from true premises to true conclusions, and so a "natural logic" should probably be concerned with them. But, and this is our point, what does Lakoff think logicians, traditional and modern, have been up to for centuries, if not the discussion and investigation of such valid, and sometimes correct, inferences?

To be precise, does Lakoff present any valid or correct inferences in his paper, as part of a proposed natural logic, that have not been extensively discussed by logicians in the normal course of their job? We would think not, and this leaves us puzzled as to what Lakoff intends the distinctive contribution of his natural logic to be.

Now there are indeed inferences to be found in Lakoff's paper, that are real world inferences, but would not be found in a logic book. However, they also have the drawback mentioned earlier, that they are not valid, or even correct, in the sense defined above.

Lakoff writes [ibid. ]:

"(34)a. Nixon refused to try to shut Agnew up. . . . (34a) entails (35a). . . (35)a. Nixon didn't try to shut Agnew up".

If Lakoff is using "entail" in its normal sense to cover valid inferences, those where the consequent must be true if the antecedent is, then what he claims is just not so. To refuse to do something is to decline, to perform a verbal act, and is so described in both American and British dictionaries. It is perfectly possible to refuse to do x and then do it, even though as a matter of fact it may be usual not to do x once you've refused to.

Again [ibid. ], Lakoff argues at length that the sentence "One more beer, and I'll leave" is derived from a sentence containing "if" such as "If I drink one more beer then I'll leave", and the intended force of the example is to show a relation of consequence between the two sentences in the derivation (of one from the other), in which case Lakoff is saying that "If I have one more beer then I'll leave" entails "One more beer and I'll leave". But that is not so, for one might neither have another beer nor leave, in which case "If I have another beer then I'll leave" is still true, but "I'll have one more beer and I'll leave" (a natural meaning of the consequent) is false, and so there can be no entailment, since the antecedent with "if" is true, and the consequent is false.

Now, we may have interpreted the whole notion of GS wrongly in that the derivation relation here is not intended to be

consequential. But IF IT IS then here again is a very shaky form of inference at the heart of the GS system: one will just not fit into the standard logical or linguistic derivational paradigms because it is necessarily making truth claims about the course of events in the real world. Such inferences can fit only into a paradigm that has the capacity to find out that it has inferred wrongly and to try again.

Note that we are not saying for a moment that we are shedding any light on difficult notions, like entailment, but only pointing out that they are difficult and unclear, have vexed logicians and philosophers, and are not nice clean tools that Lakoff, or any other linguist, can just pick up and get to work with. They need a lot of conceptual cleaning up themselves, and Lakoff shows no sign of being prepared to do that.

Lakoff's failure to provide any sort of system of rules, however miniaturised in scope, is an important one, as we argued earlier. For it leaves an important doubt as to just what a natural logic, or indeed a generative semantics, is intended to accomplish with regard to some body of sentences in a natural language. And it is not possible for Lakoff to take refuge here in the competence-performance distinction and to say that of course he is not attempting to model a speaker's performance etc. etc. , precisely because that is not what he is being accused of. As we

shall argue in section 6 below, the request for determinateness and precision is in no way to be confused with a demand for psychological explanation.

It is perfectly true that logicians import structures into their work and inform their readers that those structures represent certain natural language sentences, without ever giving a hint of a determinate translation procedure that would take us from the sentences to the structures. But we do not think that Lakoff, or any other linguist, could take shelter with the logicians here, for there is an important difference between the logician's enterprise and the linguist's. The logician is concerned above all with the formal relations between the structures he derives: the exact relation, between the structures and the natural language they "hook onto", is secondary, even though vitally important. But Lakoff, on the other hand, describes his task in terms of the production or generation of sentences along with their structures. So, for him, the missing determinateness is, and must be central.

IV. What then is a reasonable field of endeavor for linguistics? We would claim that the study of meaning is vitally important but that meaning must be studied in a new light, namely with respect to the actual usage of speakers.

Enormous strides were made in linguistics when theorists realized that the methods devised for handling phonological and

morphological data were not necessarily the best way of studying syntax. A similar mistake is presently being made by linguists studying semantics. Methods devised for studying syntax cloud the issue more than they aid it. Consider for example, McCawley's [9] sentence 'My buxom neighbor is the father of two'. McCawley considers the problem of how not to generate this sentence, or how to mark it as odd.

We submit these are two different problems, neither of which is helped by having to use the old syntactic notions to solve the semantic problem. As a problem for generation, it was perhaps reasonable to inquire how to not generate sentences that are 'syntactically bad'. Syntacticians extended the question of grammaticality to include the blocking of the generation of sentences that were 'semantically bad'. But the latter is not really a problem at all. If we are trying to account for the ability of humans engaged in the same process, we must recognize that humans generate thoughts that are meaningful within their own systems. A genuine generative system that concerned itself with generating semantically correct sentences would be doing one of two things. Either the generation would be in response to some input, (i. e. a question or statement by another person), in which case the semantics of concepts being used would already be included; or the generation would be in response to an

internal input and would utilize the data base memory of the speaker which presumably (if the person is normal) would have relationships between the data that were in accord with the speaker's conceptual experience (or 'semantics'). That is, rules for generation of sentences in a competent speaker of a language are dependent on an input that is already semantically correct. It is simply unnnecessary to worry about blocking the generation of semantically deviant sentences. Semantically deviant sentences are only generated by a speaker if they are being used to make some meaningful statement in an unusual way. No model of generation that addresses the problem of the blocking of semantically anomalous or ambiguous sentences can be seriously considered as either of model of competence or performance.

Another problem that generative linguists address is the marking of a sentence such as this as odd. Here again, from the point of view of performance, this problem makes no sense. But, within an analytic framework it is a problem and, as we have seen, generative semanticists are really doing analysis, which is why they consider the sentence a problem. An analytic system must be able to recognize this sentence as odd. But then it must do something else. It must interpret it anyway. Here then, we can have a notion of an interpretive semantics. But this interpretive semantics must act as a true interpreter. That is, it must render an apparently

anomalous sentence meaningful(by metaphor, modification of conceptual experience or whatever). This is an ability that every competent speaker has. He has it for the simple reason that nearly 100% of what he hears is meaningful, even if parts of it are in violation of certain selectional restrictions. What he has is an ability to operate on violations of selectional restrictions in such a way as to solve the problem of 'what did he mean by that?' This ability is by no means restricted to superficially anomalous utterances. Sentences such as 'Fire' must be interpreted in order to understand what to do in a dangerous situation.

What we have been adumbrating here is often made light of by linguists by classing it as a "performance" theory. Let us now look again at to what this elusive competence-performance distinction is really all about.

V. Chomsky's 'fundamental distinction' between 'competence' and 'performance' is fundamental only insofar as one wants to develop a competence grammar in the first place, and doing that is cetainly not the task we have set ourselves. Although Chomsky may have done an adequate job of providing the basis of a competence grammar, the question arises as to what the point of such a grammar is. Chomsky states that linguistic theory is mentalistic in that it is concerned with discovering a mental reality underlying actual behaviour [2]. However, results have

been largely negative when attempts to prove the psychological validity of this competence grammar have been made. Fodor and Garrett [4] comment:

> "What is one to make of such negative findings? The simplest move would be to deny the validity of the experimental procedures. . . If one is to deny the validity of such procedures in cases where they appear to fail, it seems one will equally have to deny their validity in the cases where they appear to fail it seems one will equally have to deny their validity in the cases where they appear to succeed. It is in any event now conceivable that enough negative data will eventually accumulate to make one wonder whether it is the theory that is at fault rather than the experiments. . . . . it is a mistake to claim psychological reality for the operations whereby grammars generate structural descriptions. "

Thus, some transformationalists recognize that their theory does not make any verifiable psychological claims. Now, it is at this point that many will bring up the competence-performance distinction and say:" but of course such grammars make no such claims. If you think they do, you can only have misunderstood the competence-performance distinction. For only a performance theory

could make such claims, and we do not advocate the construction of such theories at this time. We advocate the construction of only competence theories, and they come with no such "naive mechanistic" claims attached". We questioners would then be referred a remark of Chomsky's such as : "To avoid what has been a continuing misunderstanding, it is perhaps worthwhile to reiterate that a generative grammar is not a model for a speaker or a hearer. " But the matter is not so simple, for Chomsky himself writes in exacly this "naive mechanistic" mode when discussing theories. He writes of linguistic theories as making psychological claims, but does not specifically qualify what he writes so as to apply only to performance theories. A particularly revealing example is the following:

> Obviously, every speaker of a language has mastered and internalized the generative grammar that expresses his knouleege of his language. This is not to say that he is aware of the rules of the grammar or even that he can become aware of them, or that his statements about his intuitive knouledge of his language are necessarily accurate. [2]

Again, when Chomsky criticizes, for example, Yngve's phrase structure grammar on the grounds that, whether or not it can generate sentences adquately, it could never be a production model for

speakers [2], then it seems clear that he is again talking in the 'naive-mechanistic' mode.

But Chomsky cannot talk in this mode and take the competence-performance distinction at its face value. For, in a sense, the distinction was created precisely to exclude this mode of talk. We do not see these examples as mere slips of the pen by Chomsky, but take them as indicative of a deep unease about the distinction itself. A cynic might say that the real function of the distinction in current linguistics is to protect linguistic theories, all called "competence" theories of course, from any suggestion of empirical test. This is quite apparent when Chomsky defines what is to be data for a competence theory: it is, by definition, to be grammatical data [1]. Hence, of course, the whole process is circular: a competence theory cannot be tested because it is defined only with respect to data that already confirms it. This adds to the difficulties we pointed out, in Section 1 above, of viewing Chomsky's theories as scientific theories in any ordinary sense of the word "scientific".

It is particularly important for us to make this point about the enormous overuse of the competence-performance distinction in modern linguistics, because we are advocating greater empiricism in linguistic theories. The overuse we most object to is the dismissal, by Chomskyans, of any theory oriented to tests,

simulation, and real language data, as "mere performance". As we have shown , the distinction itself is simply not firm or clear enough to be used to dismiss anything whatever from consideration. The criteria of judgement in lingauistics MUST be those common to other sciences and intellectual disciplines: they cannot be simply created by fiat to protect contemporary orthodoxy.

There is some importance in pointing out, as Chomsky did in his early discussion of the competence-performance distinction, that there is no need for a linguistic theory to take account of the memory limitations, inattentions, and distractions of actual speakers. This is of course correct, but then real billiard balls when they roll pick up dust and fail to conform to the expected "competence" and performance of ideal billiard balls. Yet we do not speak of the competence and performance of billiard balls; nor is it clear that we need to introduce into linguistics a distinction unknown and unneeded in other branches of science.

The structure of Chomsky's talk about 'competence' is highly reminiscent of disputes about what it is to 'have a concept', particularly in regard to the well-canvassed philosophical possibility that a man might have the concept red, say, and never succeed in correctly picking out red stamps from a pile of colour-assorted ones. The arguments about this situation are rather like Chomsky's defence of the notion of an intrinsic competence, or

grammatical ability, if made in the case of a man who always split his infinitives. The parallel can be seen most clearly in Chomsky's discussion of the acquisition of grammar and the degree to which this requires 'specific innate abilities' and 'formal universals' [2]. If our argument has been correct then the cash-value of the notion of 'competence' simply doesn't warrant all this investment in its philosophic defence. Belief in 'competence' cannot be refuted, in any strong sense, but the question arises 'do we need to go on about it?'

Hence we claim, that , ultimately, there can only be performance models, and that when Chomsky talks of competence models he is necessarily talking about models for certain selections from among possible performances.

So then, we have argued that Chomsky's distinction is a contrived one at best; and certainly not one with sufficient power or intrinsic clarity to dismiss serious new proposals in linguistics unread. If we must place the systems we advocate , in terms of the distinction, treating it for the purpose as no more than a heuristic division , we would describe what we propose as "simulative performance. "

There is a difference separating the simulation of knowledge and linguistic processes from the modelling of actual verbal behavior. Of the former we can speak, as Chomsky does, of the ideal

speaker-hearer. Clearly the ideal speaker-hearer is not inattentive or distracted. He does, however, have memory limitations and non-linguistic knowledge. This certainly must be included as part of linguistic theory. The kind of theory of 'performance' of which Chomsky speaks may well be in the far distant future to which Chomsky relegates it. However, a theory of the kind we have been discussing is not far off. One could argue that the construction of a linguistic theory, that both accounts for the data , and does it in such a way consonant with the human method for doing so, is not all that remote. Clearly, such a theory must deal with non-linguistic knowledge and problems of human memory as well as the problems that Chomsky relegates to 'competence. In particular, after elimination of problems such as distraction, we can expect to find a linguistic theory that is neither one of 'competence' nor 'performance', but something in between and partially inclusive of both.

Chomsky writes [2]:

> The grammar does not, in itself, provide any sensible procedure for finding the deep structure of a given sentence, or for producing a given sentence, just as it provides no sensible procedure for finding a paraphrase to a given sentence. It merely defines these tasks in a precise way, A performance model must certainly incorporate a grammar; it is not to be confused with a grammar.

Thus it would be wise to take the notion of a realizable performance model as being somewhere between Chomsky' notion of competence and performance. Thus simulative performance is protected against the distractions and inattentions of real speech in a way that Chomsky originally wanted, and which seems eminently sensible. What it is not protected against is comparison with sentences produced by systematic application of a body of rules in the way that all too many "competence" theories now seem to be. The effect of this use of "competence" has been to make such work irrefutable but ultimately pointless.

VI. We have set out what we feel is wrong with the current thrust of linguistic theory, and it is perhaps time to make explicit what we feel a linguistic theory should do: i). A linguistic theory must provide rules equivalent to a mechanism for the acceptance and interpretation of normal as well as supposedly anomalous sentences.

As we pointed out above, the central discussions in current linguistic theory deal with the problem of blocking the generation of so called "starred" sentences (see Lindsay [8] for a good discussion of the ambiguity of the use of the asterisk in linguistics). We claim that linguistics must concern itself more with the interpretation, rather than the rejection, of odd sentences. To return to McCawley's example:

(1) My buxom neighbor is ,the father of two.

We claim that explicit procedures must be developed to explain why most hearers could understand that 'is the father of' here means 'acts like a father to', rather than to mark this sentence as anomalous.

Likewise, sentences (2) and (3) must be analyzed by an explicit procedure that would not refuse to accept them, but would either change its own knowledge about possible events, of add the information that something odd had happened, or that the speaker was crazy.

(2) John ate a book.

(3) I saw an elephant walk down Broadway.

ii). A linguistic theory must encompass a procedure for making explicit the information that is implicit in certain sentences.

In order to meet this requirement, we throw ourselves open to the problem of being mistaken on occasion. It is our claim that this is not unreasonable in a theory of this kind, that has recovery after failure, and learning, capabilities. Specifically, we are suggesting that sentence (4) refers implicitly to transfer of of possession and transfer of location of its object.

(4) Fred wants a book.

Furthermore it should be possible to glean from a semantic representation of (4) that it is a possible inference that (5) is true.

(5) Fred intends to read a book.

Likewise sentence (6) implies eating, both because of the use of 'have' and the use of 'dinner'.

(6) I had a steak for dinner.

Similarly a semantic analysis must include the possibility that some unknown action is being referred to in (7) and that this action is quite likely 'cooking' or something of that sort.

(7) Have you started the chicken yet?

Certainly such an analysis can be wrong. (it is easy enough to think of an alternative analysis for (7) although it would be highly involved. ) But what we are claiming is that this possibility of making a mistake is sorely needed in linguistic theories.

iii). An analysis procedure provided by a linguistic theory should proceed in such a fashion as to make what later turn out to be mistakes, when that is warranted.

As examples of this we have the above sentences and also certain syntactically ambiguous sentences such as (8).

(8) I saw the Grand Canyon flying to New York.

This sentence is ambiguous but an effective analysis procedure cannot discover both meanings at once. Rather, in this case, a good analysis procedure would, on finding the incorrect analysis first, decide that it disagreed with its semantic

information about what locations can do, and go back and modify its interpretation. That is, a linguistic theory should provide a theory of backtracking in analysis that is used to choose among alternatives.

iv). A linguistic theory must break down word meanings in a regular fashion so as to explicate the underlying elements that are in common between them.

Work on point iv has been undertaken by some researchers recently and we commend this effort. In particular Miller [10] and also Lakoff [7] have tackled this problem.

Examples that illustrate this problem are (9) and (10).

(9) John asked Mary to hit Bill.

(10) John advised Mary to hit Bill.

These sentences are very similar. A good semantic theory must point out that the elements of communication and hitting are both present in (9) and (10) and that the basic difference between the sentences lies in the implication that John believes that John will derive benefit from the hitting in (9) and that Mary will derive benefit from the hitting in (10).

v). Sentences that are identical in meaning should have identical semantic representations, and those that are similar should have similar representation.

This can best be illustrated by sentences (11) and (12) which use quite different words but basically mean the same thing.

(11) John prevented Mary from leaving the room by locking the door.

(12) Mary couldn't leave the room because John locked the door.

A good linguistic theory must explain why the concept of prevention can be referenced even in the absence of the word 'prevent', or else how the word 'prevent' refers to a combination of more basic concepts.

Similarly the old 'buy-sell' controversy about which element is more basic (i. e. see Katz [6]) is not to the point. For both sentences (13) and (14) refer to the change of possession of a book and the change of possession of money (using point II about implicit information).

(13) John bought a book from Mary.

(14) Mary sold a book to John.

A linguistic theory need mark only focus differences on identical semantic representations, if that is all that is called for (as is the case here).

vi). A linguistic theory must account for metaphor in a non-ad hoc way.

Consider sentences (15) and (16):

(15) John saw Mary's point.

(16) Bill hit upon the idea at work.

These sentences both use metaphors that are easily interpreted by analysis procedures. Each marks a change in the possible object of the verb from physical to mental and produces a concomitant change in the meaning of the verb. For example, if 'see' is transfer of physical information to a mental being, then 'see' in (15) could be the transfer of mental information to a mental being. Likewise if 'hit' requires contact with a physical object then 'hit' in (16) could be requiring contact with a mental object.

It should be clear that point vi. is directly dependent upon point iv.

vii). A linguistic theory must account for sentences in a given context rather than in isolation.

Almost all contemporary linguistic papers (including this one so far), deal with sentences only in isolation. The fact is that sentences actually occur in contexts, and the sentences often mean different things because of the contexts. While others have pointed this out before, we feel that it is important to do so again because of the inference problem. Consider sentence (17).

(17) Queen Elizabeth I had red hair.

Under all circumstances , this is a statement about Queen Elizabeth's hair color. But, depending on the sentence that preceded it, additional statements are possibly being made implicitly. If, for example, (18) had preceded (17),

(18) All red heads are mean.

then one of two possible additional statements are being made. Either the speaker, in responding to (18) with (17), is saying that Queen Elizabeth I was mean, or he is stating that (18) is false because as we all know Queen Elizabeth I was quite pleasant. Which one of these statements is actually being cannot be determined without regard to supposed common memory structures between the speakers. The important point is that a sentence can in fact have meaning apart from its own meaning structure that is derived from its opposition to , or elaboration upon, some previous sentence.

viii). A linguistic theory must lead to some precise and explicit body of analytic rules, and preferably on from there to a body of equally precise and explicit generation rules.

We argued earlier that linguistics has moved in the last fifteen years from a goal of sentence production, for which bodies of rules were sometimes written but rarely operated, to the practice of ad hoc sentence-by-sentence analysis, for which bodies of of rules are no longer thought necessary. Empiricism in linguistics has been on the wane, to put it mildly.

It is simply a fact of academic observation that the descriptions linguists provide for utterances are disputable, and disputed. The production, or non-production, of strings, by rules expressed as an algorithm provides an indisputable justification for

whatever linguistic classification and description-by-rule was initially imposed and programmed. The linguistic case is quite different from logic; for it is not usually necessary to operate a logical system very far in order to see whether or not it produces the appropriate set of strings, the theorems, for that can usually be seen by inspection. But the rules of the linguists are generally so much more numerous and complicated that inspection is not sufficient. Furthermore, inspection in such cases is prey to the well-known weakness of investigators of seeking what supports their case and ignoring what does not. If the strings are produced by algorithm, possibly out of a machine, it is more difficult to select unconsciously in that way. However, a body of analysis rules only , however precise, still leaves us with only the inscrutable structure strings. There is no doubt they were produced, but the question would remain as to what they were. However, even that situation would be some advance on the present one, where much work is merely programmatic towards the production of such algorithms.

A much stronger test situation arises if the strings produced are themselves at the surface level, after the addition of a body of generative rules. We would argue that therefore machine-translation or -paraphrase remains, in some sense, the raison d'etre of modern linguistics. For only within those enterprises can there be any real test of the vast body of work in linguistics in the last fifteen

years. Someone may argue at this point that the proposed explication would, in some sense, reduce linguistics from a science to a form of engineering. The criticism is basically correct, yet its consequences are not as are not as fearsome as some might suppose. The notion of engineering does not, of course, exclude theory: Bridges built without an adequate theory of materials simply fall down. The early attempts at machine translation failed, as any linguist would point out, because they lacked any adequate theory of linguistic structures. They were, in a sense, mere engineering. Any serious machine translation requires a classificatory theory of the algorithms to be employed. It would not be altogether wrong, historically, to say that transformational linguistics was an attempt to provide such a theory of algorithms in response to the MT debacle of the fifties, even if, as now appears likely, it was not an adequate response.

The eight points above are really only a starting place for linguistic theory. Certainly many more points could be added. We would like to point out that we are not merely presenting problems here, for we have also attempted to find solutions. Both authors have independent computer systems running at Stanford University (see [13] and [19]) that do satisfy at least some of these points. We wish here merely to point out that we feel this is the correct direction for linguistic theory to take.

VII. From much of what has been written here it should be clear that we do not think conventional linguistic theories are going to reach the goals for linguistic theory we set out in section 6. And that goes for both TG and GS. Yet it may be worth making quite clear why that is so, and in terms of a new argument. We argued for an "understanding system", one that tries to understand and interpret input sentences in context, rather than assign them to one of two heaps, the acceptable and the unacceptable, in the way that all conventional linguistic systems do, or rather, would do if they were really designed and run. Conventional linguistic systems have to act in this way: it is part of their metamathematical heritage, in which a language L is by definition a set of "acceptable" sentences. Yet, we would argue that a natural language cannot be viewed usefully as a set of sentences in any sense of those words. The reason for this, stated briefly and without the detailed treatment of [16] and [17] is that for no sequence of words can we know that it cannot be included in the supposed set of meaningful sentences that make up a natural language.

This fact, if it is a fact as we claim, has disastrous consequences for the metamathematical view of natural language as a whole, for it follows that what one might call an understanding system, an operating system of rules that was prepared, in principle,

to analyze and interpret any input, could only be represented in metamathematical terms by a self-contradictory system of rules; since, in any conventional Tarskian axiomatization, from a self-contradictory set of axioms anything whatever can be deduced, [19] and any set of "axioms" from which a randomly chosen sentence can be deduced/produced must be itself a self-contradictory system of "axioms". However, given that human beings do operate with their languages in the way described, in that their main effort is to understand and interpret whatever superficially unpromising input they receive, rather than to reject it, it seems clear that the proper deduction from the last paragraph is that it is the metamathematical analogy for language that must give way, rather than the facts of language use.

This point is closely related to another that has surfaced informally in the course of this paper. We argued the need for a linguistic theory to be able to make (possibly mistaken) inferences. By that we intend to refer to the whole area of inferences that humans make on the basis of what they see, hear, know, and remember, but which are not VALID inferences, in that they may well turn out to be wrong. For example, if we hear someone say "Please sit down", we may infer, as a matter of social habit, such things as that there is a chair in the presence of the speaker; that whatever is spoken to is human; that, in obeying the request, if he does so, the hearer will

move downwards (though he may already be lying down). Any or all of these inferences may be true, and may moreover be usually true, but may also be false on any particular occasion. These inferences are all inductive, habitual, empirical, but have no interesting logical content, because they are not valid inferences.

Our view is that such inductive rules can only be a useful part of a mechanism which is able to FOLLOW UP these, possibly mistaken, inferences to see whether or not they are justified by the information reaching the system later, and hence is also able to abandon erroneous inference where possible.

It was pointed out above, in connection with Lakoff's work, that he does make use of such inductive inferences in his informal analyses all the time, but he is mistaken if he thinks he can do that and still stay within the overall derivational paradigm of TG. For a conventional derivation, TG or GS, cannot be "run again" if it makes an error, as it surely must if it uses inductive inferences.

Work is actively proceeding on the construction of systems that can do this, by the present authors among others (see [14] and [19]. One main constraint on the form of their algorithms, is that their sub-algorithms are hierarchically organized, so that the derivations at lower levels can be rejected if necessary. [The abstract form of one such system is given in [16]]. This is never possible within any one-level system such as a "body of

transformational rules", which can reject only proferred sentences, but never its own "successful" derivations.

Multi-levelled systems of the sort we advocate belong within the discipline usually referred to as "artificial intelligence", and the goals for linguistics we advocate would draw it inevitably in that direction, and away from the Tarski-Post metamathematical paradigm of TG and GS, which has had a good run for its money but is due for a rest.

It should be pointed out that one effect of this change of direction would be to bring the subject, in some sense, back to its traditional interests. Halliday has pointed out [5] that much of what used to be called linguistics is now relegated to "sociolinguistics": the study of the relation of utterances to the physical contexts of their use and so on. There is no place for such things within the derivational paradigm, as we pointed out at length: for the modern compleat linguist utterances are simply right or wrong, as they stand and in isolation from every thing else.

The goals for linguistics we advocate, including the study of inference within and from context, and the hierarchical, intelligent, formal systems we advocate for explicating them, would bring these traditional interests back to the center of linguistics.